

Landing Site Selection using Fuzzy Rule-Based Reasoning

Navid Serrano and Homayoun Seraji

Abstract— In this paper, multiple on-board sensors are used to assess the terrain safety in real-time during spacecraft descent. A linguistic, fuzzy rule-based reasoning engine is used to determine terrain safety from sensor measurements and, together with information about required fuel consumption and site science return, provide a figure-of-merit for all possible landing sites. Additional fuzzy rule-sets are used to address spatial and temporal dependence in the reasoning process in order to arrive at a final score for each potential landing site. This landing score is used to retarget the spacecraft if the original landing site is found to be hazardous. Simulation studies are presented for illustration.

I. INTRODUCTION

ON July 20 1969, as the Apollo 11 Lunar Module was approaching the surface of the Moon, the flight commander Neil Armstrong realized that the flight computer was taking the spacecraft to a field of boulders near a crater. He immediately switched to manual control, pitched the spacecraft, cleared the hazardous terrain, and landed on a soft sandy patch. Thus, Neil Armstrong effectively saved the mission and paved the way for mankind's first steps on the Moon.

Over the past few years, there has been a concerted effort to develop technologies that enable *autonomous* safe landing of a spacecraft on potentially hazardous terrain. The autonomous landing systems have mostly been contemplated for space exploration missions to distant planets where manned flight remains out of the question. In particular, under funding from NASA/NRA, the authors have developed an autonomous multi-sensor system for safe landing on Mars called *SmartLand* [1].

The focus of this task is on the use of reasoning techniques for landing site selection during autonomous spacecraft descent. The reasoning engine must effectively integrate both on-line and off-line information that is crucial to a successful landing. Among the factors that contribute to a successful landing are terrain safety, fuel consumption, and scientific return. Terrain safety and fuel consumption are determined on-line during spacecraft descent. Multiple on-board sensors provide measurements of the terrain that are used to detect

landing hazards and assess safety. While landing safety is a critical factor in any mission, there must also be scientific value. Hence, sites with a high potential for scientific return are determined off-line by scientists prior to the mission during a rigorous and extended process [2].

A variety of techniques can be used for the reasoning process. For instance, landing site selection using probabilistic reasoning is described in [3]. In this paper, we discuss the use of fuzzy reasoning—specifically, a rule-based approach. The linguistic fuzzy rule-set, which models the expert's decision-making process, integrates terrain safety, fuel consumption, and scientific return in order to determine landing site quality. Using defuzzification, a numerical figure-of-merit (or landing score) is derived for all possible landing sites independently. In a subsequent step, spatial and temporal dependence are addressed using further fuzzy reasoning that considers landing scores earlier in the descent (temporal dependence) and from neighboring sites (spatial dependence). The final landing site is selected based on the maximum landing score after integrating all the necessary information. This fuzzy approach to landing site selection is evaluated by simulating spacecraft descent onto a variety of planetary terrains.

The paper is structured as follows. Sections II through IV discuss the key criteria for landing success; namely, terrain safety, fuel consumption, and scientific return. Section V reviews fuzzy reasoning and introduces the rule-based method used to determine landing site quality. Illustrative simulation studies are presented in Section VI. The paper is concluded in Section VII.

II. TERRAIN SAFETY

Terrain safety assessment is performed by the spacecraft's on-board sensors. Because safety is of paramount importance, it is critical to ensure that the sensors can be relied upon for terrain characterization. The use of multiple on-board sensors is thus proposed in order to provide added robustness. The three sensors considered here are a phased-array terrain RADAR, a descent camera, and a scanning LIDAR. More detail on the particulars of each sensor is provided in [1].

Using a combination of active and passive sensors with different physical characteristics such as field-of-view, resolution, and operating range provides robustness. In addition, the different operating range of each sensor leads to a tiered approach [1]. In the tiered approach, the sensors are grouped based on their ranges of operation, as shown in Table I.

Manuscript received September 15, 2006. The research described in this publication was carried out at the Jet Propulsion Laboratory, California Institute of Technology under contract from the National Aeronautics and Space Administration (NASA) with funding from the Mars Technology Program, NASA Science Mission Directorate.

N. Serrano and H. Seraji are with the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109 USA (phone: 818-393-0627; fax: 818-393-5007; e-mail: firstname.lastname@jpl.nasa.gov).

TABLE I
TIERED SENSOR OPERATION

Tier	Range	Operational Sensor(s)
1	10km – 8km	RADAR
2	8km – 1km	RADAR + Camera
3	1km - Touchdown	RADAR + Camera + LIDAR

In each descent tier, terrain features are extracted from the operational sensors. Slope and roughness features are obtained from the RADAR and LIDAR. The Least Median of Squares (LMedSq) regression technique [4] is used to locally fit a plane to the range data. Given the plane model $z = ax + by + c$ in \mathcal{R}^3 , the local slope is obtained by

$$f_{\theta}(x, y) = \cos^{-1}\left(\frac{1}{\sqrt{a^2 + b^2 + 1}}\right), \quad (1)$$

where a and b are the parameters of the best fitting plane at location (x, y) . The fitting error between the plane and the range is used as a measure of roughness:

$$f_e(x, y) = |d(x, y) - (ax + by + c)|, \quad (2)$$

where, $d(x, y)$ is range data from either the RADAR or LIDAR. Robust and computationally efficient hazard detection algorithms are used to locate craters and rocks from camera imagery. Crater boundaries are represented by an ellipse [5]. Let $x_{0,i}$, $y_{0,i}$, a_i , b_i , and ϕ_i be the ellipse center x -coordinate, center y -coordinate, semi-major axis length, semi-minor axis length, and rotation angle, respectively, for the i th detected crater. The crater map is then

$$f_c(x, y) = \begin{cases} 1, & \text{for } \left(\frac{x^2}{a_i^2} + \frac{y^2}{b_i^2}\right) \leq 1, \\ 0, & \text{otherwise} \end{cases}, \quad (3)$$

where x and y are points in a coordinate system rotated by ϕ_i and translated by $x_{0,i}$ and $y_{0,i}$.

At lower altitudes, rocks and boulders are visible and are detected using the algorithm described in [6]. Rock sizes are estimated based on shadow projection patterns and the known sun angle. The rock map is simply

$$f_r(x, y) = \begin{cases} 1, & \text{for } (x, y) \in R \\ 0, & \text{otherwise} \end{cases}, \quad (4)$$

where R is the set of pixel locations in the image identified as rocks. Example detection results are shown in Figure 1.

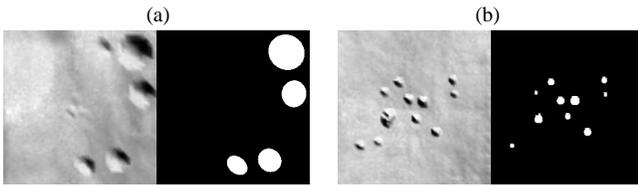


Fig. 1 Craters (a) and rocks (b) detected from camera imagery.

III. FUEL CONSUMPTION

As the terrain safety is assessed during descent, retargeting operations can be performed in order to avoid landing hazards and reach a safe landing site. However, the reachable terrain is constrained by the spacecraft's descent trajectory, velocity and available fuel. Using ballistic analysis, it is shown in [7] that the reachable terrain (landing footprint) is bounded by the fuel ellipse

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1, \quad (5)$$

where the semi-major axis a and semi-minor axis b are computed from

$$a = \frac{(\Delta V^2 - 2E/m)\Delta t}{2\Delta V(1 - v_H/\Delta V)}; \quad b = a\sqrt{1 - v_H/\Delta V}. \quad (6)$$

In the above equations, ΔV is the allowable change in velocity based on fuel allocation, Δt is the time to impact, v_H is the horizontal velocity, m is spacecraft mass, and E is kinetic energy [7].

Figure 2 shows estimated landing footprints at various points during descent. The position of the spacecraft is shown as a red circle. The ballistic trajectory begins with an initial horizontal velocity. The landing ellipse changes dramatically after a re-targeting maneuver is applied.

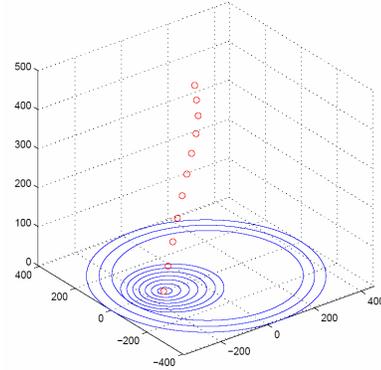


Fig. 2 Estimated landing ellipse using ballistic analysis.

IV. SCIENTIFIC RETURN

Landing site selection for a space exploration mission is generally a compromise between terrain safety and scientific return. When safety cannot be guaranteed, a potential site must be discarded—regardless of its potential scientific impact. Determining areas of high scientific return is a laborious process that involves numerous considerations beyond the scope of an on-board reasoning system [2]. It is, however, possible to integrate the scientists' preferred sites in order to influence the on-board site selection. Thus, for instance, the scientists may pre-select multiple potential sites that can be used in conjunction with the on-board terrain safety assessment in order to select the best site during descent. Such a scenario is considered here.

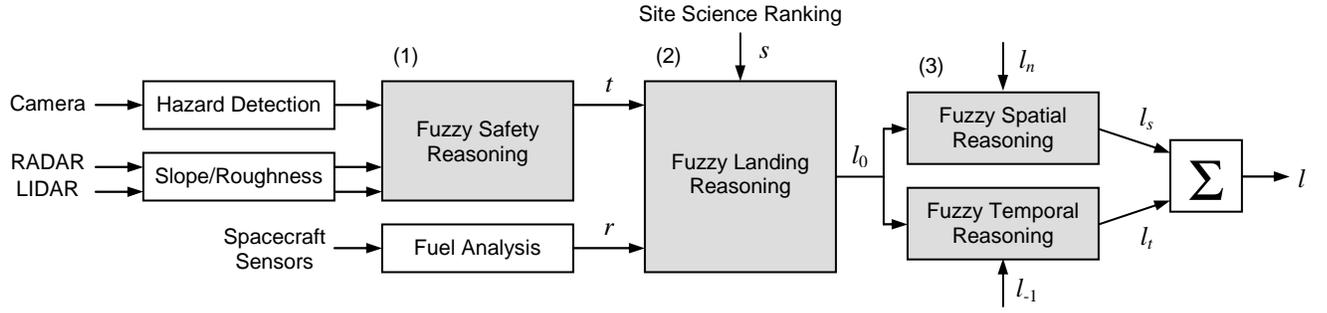


Fig. 3 Block diagram of the proposed landing site selection approach. The three phases of fuzzy reasoning are shown in gray.

Assume scientists select a set of points of interest $(x_{0,i}, y_{0,i})$ on the terrain. Each point is at the center of a circular region of interest with radius r_i :

$$(x - x_{0,i})^2 + (y - y_{0,i})^2 \leq r_i^2. \quad (7)$$

The regions of scientific interest may or may not be ranked. If the sites are ranked, the ranking may be relative to other sites or based on a scale of interest. At the point of entry, all pre-selected locations are reachable. As the terrain safety is assessed, the site that best combines terrain safety, fuel consumption, and scientific return is used for re-targeting.

V. FUZZY REASONING

The field of fuzzy logic was introduced by Zadeh in 1965 [8]. A wide variety of practical applications using fuzzy principles have been demonstrated over the years, including relevant autonomous tasks such as navigation [9] and landing [10]. Part of the appeal of fuzzy systems is that they can be used for approximate reasoning. This is particularly important when there is uncertainty in the reasoning process, in addition to imprecision in the data. In the context of this paper, the sensor measurements are the source of uncertainty, as they can be corrupted by noise. Observe that the use of *linguistic* fuzzy sets and simple IF-THEN rule statements enables fuzzy logic to model a human expert's reasoning and decision making process.

A. Rule-Based Approach

The fuzzy rule-based approach consists of a set of linguistic statements, or rules, defined by a human expert. Each rule is of the form

$$\text{IF } C, \text{ THEN } A, \quad (8)$$

where the *condition* C is composed of fuzzy input variables (e.g. terrain safety, fuel consumption, scientific return) and fuzzy connectives (e.g. AND, OR, NOT) and the *action* A is a fuzzy output variable (e.g. landing site quality). The rules are evaluated based on their membership to fuzzy sets. As opposed to traditional Boolean logic that requires full membership in a set, a fuzzy system includes a set of membership functions that allow for degrees of membership in multiple sets.

Determination of the best landing site is performed in

three fuzzy reasoning phases as depicted in Figure 3. In the first phase, the safety of the terrain is assessed using a set of rules based on features extracted from the on-board sensors, as described in Section II. The second phase involves integrating information based on terrain safety, fuel consumption, and scientific return using a landing quality rule set. In the third phase, the landing quality at earlier times in the descent and the landing quality of neighboring sites are combined using spatial and temporal rule sets.

B. Fuzzy Terrain Safety Assessment

Terrain safety, t , is represented by four fuzzy sets with the linguistic labels {P, L, M, H}, which stand for *poor*, *low*, *moderate*, and *high*, respectively, as shown in Figure 4.

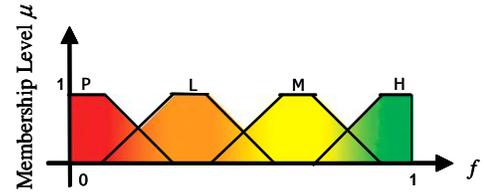


Fig. 4 Fuzzy membership classes.

The RADAR and LIDAR both yield range data that is used to extract slope and roughness features. Hence, as shown in Figure 5, the same rule-set is used for these two sensors.

The linguistic labels for the slope f_θ are {VF, F, S, VS}, which stand for *very-flat*, *flat*, *steep*, and *very-steep*, respectively. The linguistic labels for roughness f_e are {VS, S, R, VR}, which stand for *very-smooth*, *smooth*, *rough*, and *very-rough*, respectively. All rules in Figure 5 are connected via the AND operator. Thus, for instance, the first rule is IF (f_θ is VS) AND (f_e is VR), THEN (t is P).

		f_e			
		VR	R	S	VS
f_θ	VS	P	P	L	M
	S	P	L	M	M
	F	L	M	M	H
	VF	M	M	H	H

Fig. 5 Rules for terrain safety from range data.

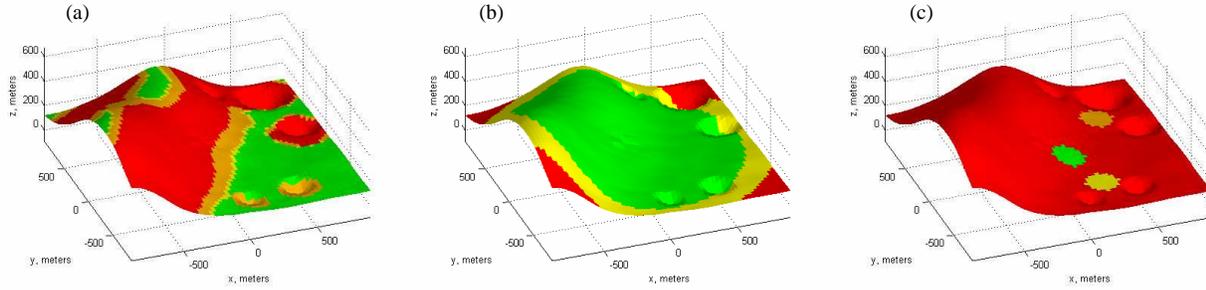


Fig. 6 Safety (a), fuel (b), and science (c) maps overlaid on example terrain.

A numerical safety score is obtained using centroid defuzzification. The safety score is a weighted combination of the degrees of membership to the fuzzy classes:

$$t_i(x, y) = \frac{\sum_j p_j A_j}{\sum_j A_j}, \quad (9)$$

where $t_i(x, y)$ is the defuzzified terrain safety score for the i th sensor at point (x, y) on the terrain, p_j is the peak value associated with the membership functions, and A_j is the area under the truncated membership function. The resulting safety score is in the range $[0.0, 1.0]$.

Let t_1 , t_2 , and t_3 represent the terrain safety scores for the RADAR, camera, and LIDAR respectively. The terrain safety score for the camera can be obtained from the hazard detection algorithms directly. In Tier 2, only craters are detected and hence, the camera terrain safety score t_2 is simply

$$t_2(x, y) = 1 - f_c(x, y). \quad (10)$$

In Tier 3, both craters and rocks are detected and the cameraterrain safety score t_2 becomes

$$t_2(x, y) = [1 - f_c(x, y)][1 - f_r(x, y)], \quad (11)$$

where f_c and f_r are the crater and rock detection maps defined earlier. Having obtained safety score for each of the sensors, a fused safety score is obtained using a weighted average:

$$t(x, y) = \sum_i \beta_i t_i(x, y), \quad (12)$$

where $\sum \beta_i = 1$. The weights β_i represent a measure of certainty associated with each sensor. The weights can be set based on environmental factors, as in [1]. The reliability of the hazard detection algorithms can also be used. For instance, if a crater is detected, that region of the terrain is known to be unsafe. However, if a crater is not detected, it cannot be said with certainty that the terrain is safe and thus, the RADAR and/or LIDAR should be weighted more. Such an approach is adopted here; camera certainty is set to zero when a hazard is not detected. Consequently, the weights β_i are also a function of location (x, y) . In Figure 6a, the fused

terrain safety score t is overlaid on a Digital Elevation Model (DEM) of a planetary landscape. The color coding is as follows: red, orange, yellow, and green correspond to P, L, M, and H, respectively.

C. Fuzzy Landing Site Quality Assessment

The three key factors for landing success (terrain safety, fuel consumption, and scientific return) can be linked to landing quality using a fuzzy rule set. Safety is derived from terrain sensors using fuzzy reasoning, as described in the previous section.

In the case of fuel consumption, a ballistic model is used to determine the elliptical boundary of the reachable terrain. However, this boundary is merely an estimate and therefore, it is important to consider a margin of error. A second more conservative boundary is introduced, thus accounting for the margin of error. Let r represent reachable terrain based on fuel consumption. Three classes are considered: {U, M, R}, which stand for *unreachable*, *marginally-reachable*, and *reachable*. Points outside the original boundary are automatically *unreachable*. Points within the new (i.e. reduced) boundary are *reachable*. And points that lie between the original and new boundary are *marginally reachable*. An example fuel map is shown in Figure 6b. The color coding is as follows: red, yellow, and green correspond to U, M, and R, respectively.

Whereas terrain safety and fuel consumption are determined on-line during descent, points of scientific interest are determined off-line, prior to mission launch. Scientists select multiple points on the terrain that are of interest. Let s represent the level of scientific return for any point (x, y) on the terrain. Four levels (or classes) of scientific return are considered: {N, L, M, H}, which stand for *none*, *low*, *medium*, and *high*, respectively. Before mission launch one of the regions is selected as the nominal landing site. Figure 6c shows three regions of scientific interest overlaid on the terrain. The color coding is as follows: red, orange, yellow, and green correspond to N, L, M, and H, respectively.

During descent, as the terrain safety is assessed, the spacecraft may retarget to another potential landing site after incorporating all relevant information. This process involves fuzzy reasoning. Let l represent landing site quality. As with

terrain safety, four classes are considered. The linguistic labels {P, L, M, H} represent *poor*, *low*, *medium*, and *high* landing quality, respectively. The landing quality membership functions are the same as those in Figure 4. The landing quality rule-set integrates the three input variables terrain safety, fuel consumption, and scientific return. The first rule addresses the worst-case scenario:

$$\text{IF } (t \text{ is P}) \text{ OR } (s \text{ is N}) \text{ OR } (r \text{ is U}), \text{ THEN } (l \text{ is P}), \quad (13)$$

where t is terrain safety, s is scientific return, r is terrain reachability, and l is landing site quality. The remaining rules employ the fuzzy connective, AND, as shown in Figure 7.

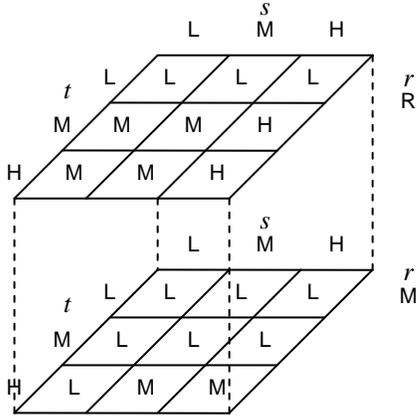


Fig. 7 Rules for landing site quality l .

D. Fuzzy Spatial and Temporal Landing Assessment

The fuzzy reasoning process described up to this point is performed for each point (x, y) on the terrain independently. Realistically, however, the quality of a landing site is neither spatially nor temporally independent. The landing score at a point on the terrain at one time during descent does not change dramatically at a subsequent time. Similarly, the landing score at a particular point on the terrain is, in general, not substantially different from the landing scores of its neighbors. Consistent with the reasoning architecture described earlier, we also address spatial and temporal dependence in a fuzzy framework.

Regional measures for safety have been addressed in prior work using filtering techniques [11]. The use of fuzzy rule-sets to incorporate spatial and temporal information can be thought of as a non-linear approach to the problem. The spatial and temporal fuzzy rule sets are shown in Figure 8.

		l_n				l_1			
		P	L	M	H	P	L	M	H
l_0	P	P	P	P	L	P	P	L	L
	L	P	L	M	M	L	L	M	M
	M	P	M	M	H	L	M	M	H
	H	L	M	H	H	M	M	H	H

Fig. 8 Rules for spatial (a) and temporal (b) landing assessment.

As shown in Figure 8, l_0 represents the landing score at the current location and current time-frame, l_1 represents the landing score at the current location and previous time-frame, and l_n is the median score of the eight cells neighboring l_0 . The spatial l_s and temporal l_t landing scores are obtained by applying the rules in Figure 8 and are then combined using a weighted average to arrive at a final landing score, l :

$$l(x, y) = \alpha_s l_s(x, y) + \alpha_t l_t(x, y), \quad (14)$$

where $\alpha_s + \alpha_t = 1$. The weights can be set equally or can be biased towards either the spatial score or the temporal score. Incorporating the spatial and temporal rules provides another layer of uncertainty management. The landing score l thus incorporates the three key landing factors (terrain safety, fuel consumption, and scientific return), as well as spatial and temporal information that mitigates spurious sensor measurements. Site selection merely involves finding the point on the terrain with the highest score l :

$$(x^*, y^*) = \arg \max_{(x, y)} \{l(x, y)\}. \quad (15)$$

The selected landing site (x^*, y^*) , which has a corresponding landing score of $l^* = \max \{l(x, y)\}$, can be used to retarget the spacecraft and ensure a successful landing.

VI. SIMULATIONS

The proposed approach is evaluated by simulating a spacecraft's descent onto a diverse set of planetary terrains closely resembling the Martian landscape. A total of ten different synthetic DEMs are used for validation. Descent onto each DEM is performed using DSENDS [12], a high-fidelity dynamics and spacecraft simulator for entry, descent and landing. RADAR, camera, and LIDAR sensor measurements of the synthetic terrains are obtained at multiple points during descent using appropriate models. Having extracted terrain features from the sensor measurements, the fuzzy reasoning engine is used to obtain landing scores for each point on the visible terrain segments.

Landing site selection results are shown in Figure 9. The safety assessment is overlaid on each terrain. The fuel ellipse is dashed and the sites of scientific interest are solid. Each site of scientific interest is centered about the original point selected by a scientist and shows a broad area with a circular radius of 100m. The science ranking is indicated by H (high), M (medium), or L (low). The sites of scientific interest are not selected by actual scientists—they are only meant for evaluation purposes. The final selected landing site is shown with a black hash mark. In addition, the final landing score l^* is shown for each selected site. All examples are at an altitude of 4km.

As can be seen from Figure 9, the fuzzy landing site selection process adequately combines the relevant factors. For instance, in Figure 8a, there are two sites with *high*

scientific return. However, one of them lies in a very unsafe portion of the terrain and the other is just beyond the reachable boundary. As a result, the site with *medium* scientific return is selected. In Figure 8b, on the other hand, the only site with *high* scientific return is selected because there is sufficient safe terrain and it is within the reachable boundary.

VII. CONCLUSIONS

This paper describes a fuzzy rule-based approach to landing site selection during autonomous spacecraft descent. Terrain safety is determined using a fuzzy rule set that integrates information from multiple on-board sensors. In addition, reachable regions of the terrain are determined using a ballistic descent trajectory based on spacecraft fuel consumption. Scientific return is also considered by allowing mission scientists to pre-select regions of interest as candidate landing sites. All three key criteria (terrain safety, fuel consumption, and scientific return) are integrated using a fuzzy rule-set. Further robustness is added by incorporating spatial and temporal information to the reasoning process. The landing site selection is performed by choosing the point on the terrain with the highest landing score after completing the reasoning process. Simulation experiments successfully demonstrated the selection of landing sites that best combine safety, fuel, and science criteria. Future work will involve more rigorous validation, including Monte Carlo simulations and experiments on real data.

REFERENCES

- [1] N. Serrano, M. Bajracharya, A. Howard, and H. Seraji, "A Novel Tiered Sensor Fusion Approach for Terrain Characterization and Safe Landing Assessment," *IEEE Aerospace Conference*, Big Sky, MT, March 2006.
- [2] M. Golombek et al., "Selection of the Final Four Landing Sites for the Mars Exploration Rovers," *Lunar and Planetary Science XXXIV*, 2003.
- [3] N. Serrano, "A Bayesian Framework for Landing Site Selection During Autonomous Spacecraft Descent," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Beijing, China, October 2006.
- [4] A. Johnson, A. Klumpp, J. Collier, and A. Wolf, "LIDAR-Based Hazard Avoidance for Safe Landing on Mars," *AAS/AIAA Space Flight Mechanics Meeting*, Santa Barbara, CA, February 2001.
- [5] Y. Cheng and A. Ansar, "Landmark Based Position Estimation for Pinpoint Landing on Mars," *IEEE International Conference on Robotics and Automation*, Barcelona, Spain, April 2005.
- [6] A. Huertas, Y. Cheng, and R. Madison, "Passive Imaging Based Multi-cue Hazard Detection for Spacecraft Safe Landing," *IEEE Aerospace Conference*, Big Sky, MT, March 2006.
- [7] S.R. Ploen, H. Seraji, and C.E. Kinney, "Determination of spacecraft landing footprint for safe planetary landing," *IEEE Transactions on Aerospace and Electronic Systems*. (To appear.)
- [8] L.A. Zadeh, "Fuzzy Sets", *Information and Control*, 12, pp.338-353, 1965.
- [9] H. Seraji: "Fuzzy Traversability Index: A New Concept for Terrain-Based Navigation", *Journal of Robotic Systems*, 17(2), pp.75-91, 2000.
- [10] A. Howard and H. Seraji, "Multi-Sensor Terrain Classification for Safe Spacecraft Landing," *IEEE Transactions on Aerospace and Electronic Systems*, 40(4), October 2004, pp.1122-1131.

- [11] H. Seraji, "Safety Measures for Terrain Classification and Safest Site Selection," *Autonomous Robots*, 21(3), pp.211-225, 2006.
- [12] J. Balaram, R. Austin, P. Banerjee, T. Bentley, D. Henriquez, B. Martin, E. McMahon, G. Sohl, "DSENDS - A High-Fidelity Dynamics and Spacecraft Simulator for Entry, Descent and Surface Landing," *IEEE Aerospace Conference*, Big Sky, MT, March, 2002.

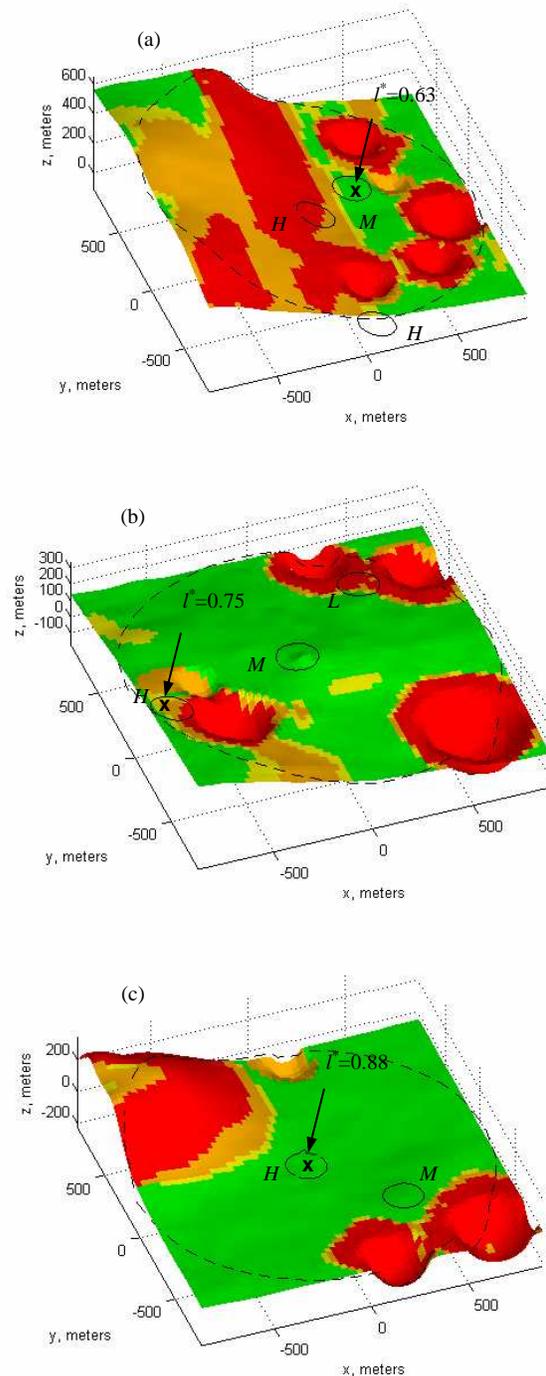


Fig. 9 Example landing site selection results for three different terrains.